

UNCLASSIFIED//FOR OFFICIAL USE ONLY

WN 98W000033

MITRE REPORT

Impact of ATM on JADS

April 1998

Devaraj Sahu

MITRE

Washington C³ Center
McLean, Virginia

Approved for public release.



UNCLASSIFIED//FOR OFFICIAL USE ONLY

20000824 040

AQI00-11-3755

Abstract

This report investigates the feasibility of an Asynchronous Transfer Mode (ATM) network to support Test and Evaluation (T&E) requirements such as the Joint Advanced Distributed Simulation (JADS) phase two end-to-end (ETE) test. It gives a broad overview of ATM and discusses how an ATM network could support JADS ETE T&E efforts.

This report reviews key concepts behind ATM such as service classes, quality of service, and available signaling types. Two methods of running Internet Protocols (IP) applications over ATM are discussed. These two methods are: classical IP over ATM and Local Area Network Emulation (LANE). Applications running ATM in the native mode are also discussed.

A JADS phase two ETE test configuration for T&E is discussed and a potential configuration for running this test over an ATM network is presented. Equipment costs and ATM service costs are given. Questions raised by the JADS T&E community are answered.

Table of Contents

Section	Page
1. INTRODUCTION	1
1.1 BACKGROUND	1
1.2 PURPOSE	1
1.3 SCOPE	2
1.4 DOCUMENT ORGANIZATION	2
2. ATM	5
2.1 ATM OVERVIEW	5
2.1.1 What is ATM?	5
2.1.2 Why ATM?	6
2.1.3 ATM Reference Model	7
2.1.4 ATM Traffic Classes	9
2.1.5 AAL Types and QoS	10
2.1.6 ATM Signaling and Addressing	12
2.2 IP and ATM	15
2.2.1 Classical IP over ATM	15
2.2.2 LAN Emulation (LANE)	16
2.2.3 MPOA	16
2.3 Native ATM Services	17
3. USING ATM TO SUPPORT ADS FOR T&E	19
3.1 JADS ETE Phase Two Test Over ATM	22
3.2 Issues and Solutions	22
3.3 Cost Analysis	25
4. CONCLUSIONS	27

Section 1

1. INTRODUCTION

The Joint Advanced Distribution System (JADS) program is chartered by the Office of the Secretary of Defense (OSD) to investigate the utility of Advanced Distributed Simulation (ADS) for Test and Evaluation (T&E) applications. JADS has also been asked to identify ADS constraints and methodologies when used for T&E and to identify requirements for ADS systems to better support T&E in the future.

The JADS Joint T&E project consists of three test programs: the System Integration Test (SIT), End-to-End (ETE) test, and finally Electronic Warfare (EW) test. The SIT test has been scheduled to run from 1995 through 1998 and the EW and ETE tests will run through 1999. For each phase of a given test program, there is an associated test scenario that defines a prescribed set of interactions among the program entities. For JADS T&E efforts, a key element of the test program is the communications infrastructure that supports the execution of a test scenario.

In anticipation of extending the application of Asynchronous Transfer Mode (ATM) technology to support simulation applications for T&E, the JADS Test Director requested that MITRE review the status of ATM technology and its capability to support T&E requirements such as those represented by SIT, ETE, and EW tests. In response, MITRE's Network and Communications Engineering Center initiated a study to address the impact of ATM technology for distributed simulation based T&E. This paper looks at one phase of the JADS ETE test program and discusses how ATM could be used to support this test and what impact it could have from a cost and performance perspective.

1.1 BACKGROUND

ATM technology is gaining acceptance in both private and government communications networks. There are three main reasons for this increased acceptance. First, the equipment vendors are making standards-compliant products that are interoperable. Second, as the ATM technology matures, prices are steadily coming down. Finally, the availability of high-bandwidth pipes supporting traffic with different Quality of Service (QoS) requirements for transport of ATM traffic makes ATM a well-liked candidate in the wide area network (WAN). The T&E community is becoming aware of the acceptance of ATM technology and wants to find out what impact ATM will have on distributed simulation.

1.2 PURPOSE

This paper has three purposes. The first purpose is to identify and document aspects of ATM that may be of interest to the T&E community for T&E applications. The second

purpose is to examine the use of ATM to support a distributed simulation exercise. The final purpose is to provide answers to issues of concern to the T&E community about ATM networks. Since cost is a real concern in fielding any network, approximate cost estimates of equipment and services are given. It is hoped that this paper will provide personnel in the T&E community information that helps them in the application of this complex technology.

1.3 SCOPE

This paper provides a high-level review of ATM. Additional and more comprehensive information about ATM can be found in the reference section.

In this paper, we examine a representative JADS ETE test that can be run over an ATM network. An evaluation of all JADS tests is beyond the scope of this paper. Finally, the cost estimates given in this paper are meant to be guidelines only and are not firm quotes. More importantly, these cost estimates will change over time.

A list of JADS ATM questions addressed by this paper follows:

- Would the ATM network be public or private? If public, what impact would other traffic have on latency of JADS data?
- What are the encryption issues? How can encryption be accomplished between sites? Would we have to invest in NES rather than KIV-7?
- What is the impact on bandwidth?
- What is the impact on latency? Will latency be fixed or variable?
- Will small packet size require more processing at simulation sites? Will this add variable latency equal to or greater than WAN savings?
- Will multicast be easier or harder with ATM?
- Will ATM be more or less reliable?
- What are cost implications? Will T1 costs go up? Will new boxes be needed at simulation sites?
- What would it take for us to build our own private ATM network (i.e. cost, equipment, training, etc.)?

1.4 DOCUMENT ORGANIZATION

Section 2 gives a broad overview of ATM. Many of the key concepts behind ATM such as service classes, quality of service, and signaling are discussed. Several options for running Internet Protocol (IP) traffic in an ATM network are also discussed. Work on running applications directly over ATM, referred to as "native mode" ATM is also discussed.

Section 3 presents a JADS ETE test scenario that can be run over an ATM network. Answers to issues of concern to the JADS T&E community are also discussed, and a cost analysis of equipment and services is given. Finally conclusions are presented in Section 4.

Section 2

2. ATM

In this section we highlight some aspects of ATM believed to be important to the T&E community. We first present a general overview of ATM. Next we discuss issues concerning IP and ATM. Finally, we discuss native ATM services.

2.1 ATM OVERVIEW

The ATM specifications are formulated by the International Telecommunications Union (ITU) and the ATM Forum. The ITU is a United Nations special agency responsible for formulating telecommunications standards. The ATM Forum is an international non-profit organization formed with the objective of accelerating the use of ATM (Asynchronous Transfer Mode) products and services through a rapid convergence of interoperability specifications. In addition, the Forum promotes industry cooperation and awareness [for more information see their web page: www.atmforum.com]:

The ATM Forum has three types of members: Principal, Auditing, and Passive. A Principal member is an active member of the Forum and pays \$10,000 annually towards membership. The Auditing and User membership are both passive with annual membership dues of \$2,000 and \$1,500 respectively. The Forum has a Technical Committee which has five meetings a year, each lasting a week. Only Principal Members may contribute papers for the working group specifications. There are twelve working groups in the Technical Committee. These working groups are: Control Signaling, LAN Emulation/Multiprotocol, Network Management, Physical Layer, Residential Broadband, Routing Addressing, Security, Service Aspects and Applications, Testing, Traffic Management, Voice and Telephony over ATM, and Wireless ATM.

The ATM Forum has been instrumental in developing and documenting specifications for ATM internetworking and promoting interoperability among vendor equipment. It also writes extensions to the ITU standards. In the following, we explain some of the key concepts behind ATM.

2.1.1 What is ATM?

ATM is a packet-switching technology that uses fixed-length packets called cells. Each cell is 53 octets long with a 5 octet header and a 48 octet payload. The payload can contain voice, data, or video information. In fact, these three kinds of traffic can be carried on the ATM network at the same time.

The asynchronous in ATM means that the cells originating from different traffic sources are *statistically multiplexed* as they arrive at an ingress ATM device such as a switch. At the

receiver end the cells are demultiplexed in the exact order they arrive at the egress switch. There is no resequencing of cells because they do not arrive out of order. The asynchronous nature of ATM traffic should be contrasted with synchronous traffic such as Digital Signal Level 3 (DS3 with a capacity of 44.74 megabits per second (Mbps)) transport which uses Time Division Multiplexing (TDM). In a synchronous transfer mode, TDM slots are synchronously transferred, whether there is traffic or not, thus wasting bandwidth. In contrast, ATM does not reserve slots for a pre-assigned traffic stream.

The logical ATM connections of a cell are determined by the Virtual Path Identifier (VPI) and Virtual Channel Identifier (VCI) fields in a cell header. The VPI/VCI combination is meaningful in the context of a given interface. The communication channels that transport ATM cells are called Virtual Connections. There are two types of virtual connections: Virtual Path Connection (VPC) and Virtual Channel Connection (VCC). A VPC is a collection of VCCs. A VCC is identified by a set of VPI/VCI values between two end points of the network. Connections can be established manually (Permanent Virtual Circuit (PVC)) or dynamically (Switched Virtual Circuits (SVC)). As the name implies, PVCs remain established permanently. In a PVC scenario, both point-to-point and point-to-multipoint connections can be established. SVCs are set up on a need-only basis using the signaling protocols of the ATM Forum. The duration of SVCs varies depending on the needs of the application. SVC makes manual setting up of connections unnecessary.

2.1.2 Why ATM?

The use of packet-switching by ATM is advantageous for several reasons. Cell-switching (a cell being a short packet) can handle both delay-sensitive traffic (such as voice and video) and loss-sensitive traffic (such as data) all at the same time. Multiple traffic streams with different traffic characteristics can share the same physical path in the network. Furthermore, ATM provides a range of high-speeds- from a few Mbps up to 10 Gigabits per second (Gbps). Finally, cell-switching can provide a broadcasting capability whereas circuit-switching can not. The benefits of using ATM are summarized below [1]:

- ATM technology is scalable.

An ATM cell can be transported over fiber or copper with a choice of line rates and framing protocols. The ATM standards simply specify the cell size, but do not mandate a line rate or a framing protocol. As a result, ATM cells can propagate over an Ethernet Local Area Network (LAN), then transit through gigabit ATM switches, and then get transmitted over DS3 lines. Thus ATM networks can serve the needs of users with a variety of equipment and a variety of bandwidth requirements.

- ATM technology is application transparent.

ATM is suitable equally for different applications such as voice, video, and data. All three kinds of traffic can exist in the same network and work seamlessly. ATM handles both delay sensitive traffic and loss sensitive traffic equally well without performance problems.

- ATM provides connection flexibility.

ATM nodes can be networked together in a variety of ways. One option may be to subscribe to a provider that sets up PVCs among the nodes. Another option may be to subscribe to SVC service that carries traffic from other subscribers. Pricing structures may be negotiated with the providers for use of bandwidth on a best-effort basis or on service guarantees,

- ATM is based on ATM Forum specifications.

There is a growing realization among equipment manufacturers that to be successful in the enterprise environment, equipment must interoperate. Network administrators feel that their investments in equipment will not be obsolete because of proprietary solutions provided by one vendor. The ATM Forum strives to promote interoperability of ATM equipment by developing specifications to promote interoperability.

2.1.3 ATM Reference Model

The ATM is based on the Broadband Integrated Services Digital Network (B-ISDN) service model. B-ISDN, in turn, relies on ISDN. The idea behind conventional (i.e. narrowband) ISDN is that a digital bit pipe exists between a customer and a carrier [2]. There are two types of digital pipes- the basic rate (BR) digital pipe and the primary rate (PR) digital pipe. Each of these pipes supports multiple channels that are interleaved by TDM. In the BR case, there are two 64 Kilobits per second (Kbps) digital pulse-code modulated (PCM) channels for voice or data (called B1 and B2 channels) and a 16 Kbps signaling channel (called the D channel) for out-of-band signaling. In the PR case, there are twenty-three B channels (each being a 64 Kbps channel) and one D channel also at 64 Kbps capacity. Since ISDN deals with 64 Kbps channels, it is called narrowband ISDN service. In contrast B-ISDN uses channel capacities much higher than 64 Kbps, namely, above 1.54 Mbps. Another important difference between ISDN and B-ISDN is that ISDN is a circuit-switched service (like a phone line), while B-ISDN is not. In other words, ISDN switches can not be used for B-ISDN switching.

The B-ISDN-based ATM reference model is shown in Figure 1. This model consists of three layers [2]: the physical layer, the ATM layer, and the ATM Adaptation Layer (AAL). The physical layer deals with electrical and optical parameters and bit timings in the physical

medium that transports the electrical signals. The physical layer consists of two sub-layers: the transmission convergence (TC) sublayer and the physical medium dependent (PMD) sublayer. The PMD sublayer provides physical interface to the cables and handles bit timings. The TC sublayer converts an incoming cell stream into bits at the ingress port and converts an incoming bit stream into a cell stream at the egress port. Thus the TC sublayer delineates cell boundaries and generates header check sums and verifies them. The next layer is the ATM layer, which provides some of the key functions of ATM [3]. The ATM layer generates and extracts ATM cell headers, provides VPI/VCI translation, and does cell multiplexing and demultiplexing. Finally, the AAL layer adapts data from higher layers into formats that conform to the ATM layer. Several different AAL data types have been defined by the ITU (e.g. AAL 1, AAL 3/4, and AAL 5). These are discussed in the following section. The AAL functions are end-to-end. The AAL layer is divided into two sublayers (Figure 2), the convergence sublayer (CS) and the segmentation and reassembly (SAR) sublayer. The SAR sublayer chops packets into cells (segmentation) at the source and puts them back into packets (reassembly) at the destination. The CS sublayer divides user data into manageable fixed-length packets called Protocol Data Units (PDUs).

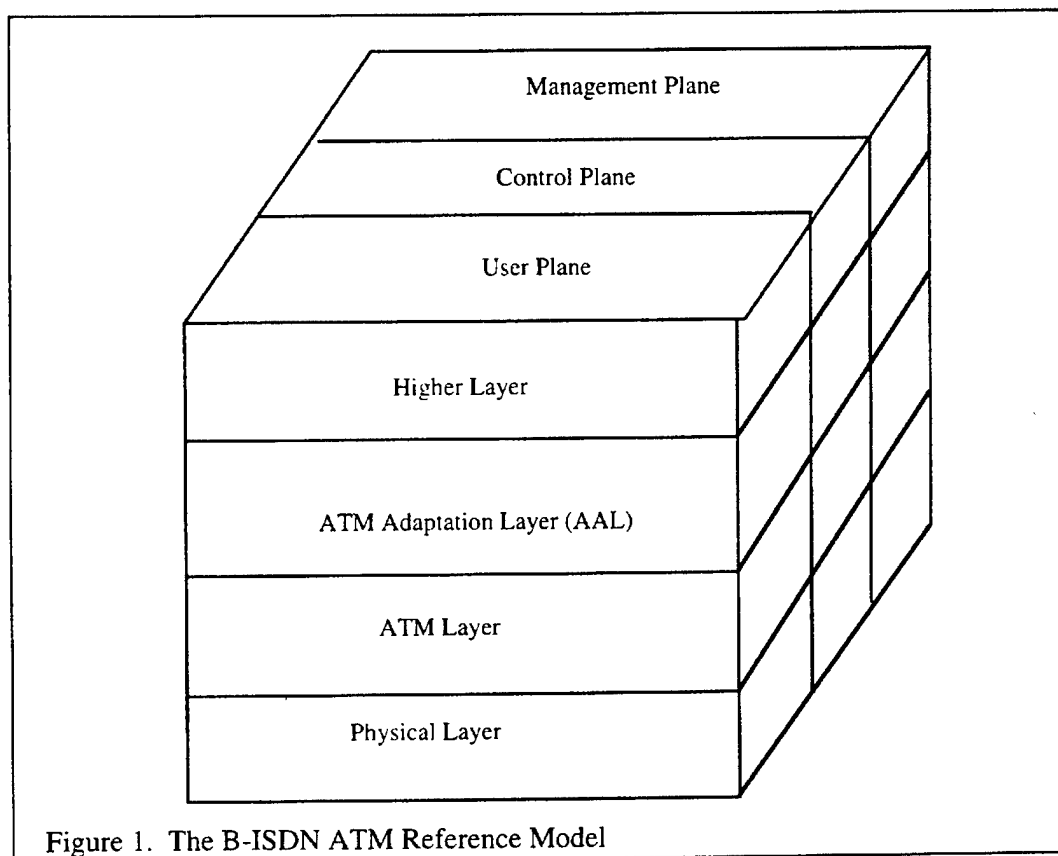



Figure 1. The B-ISDN ATM Reference Model



In addition to the layers, the reference cube consist of three planes: the user plane, the control plane, and the management plane. The user plane spans the Physical Layer, ATM Layer, AAL Layer, and higher layers [3]. The higher layers run applications which interface with the AAL Layer via application programming interfaces (APIs). The control and management planes support the services provided by the user plane. The control plane supports control functions such as signaling and connection establishment. The management plane supports layer management and plane management. Layer management monitors user and control plane for faults, generates alarms, and initiates corrective action. Plane management performs coordination across all layers and planes.

2.1.4 ATM Traffic Classes

The ATM layer is insulated from the needs of user applications. Thus there is a gap between the user needs and the services provided by the ATM layer. The AAL layer bridges this gap. As its name implies, the AAL layer performs the functions needed to adapt the services provided by the ATM layer to the services required by applications. In an ATM network there are different AAL types, which support different traffic classes. ATM traffic classes, AAL types, and ATM QoS parameters are discussed.

There are five ATM traffic classes currently defined: Constant Bit Rate (CBR), Real Time Variable Bit Rate (RT-VBR), Non-real Time VBR (NRT-VBR), Available Bit Rate (ABR), and Unspecified Bit Rate (UBR). Each is discussed further below.

The CBR traffic class transmits bits at a constant rate. A Timing relationship between the source and destination is required. This traffic class covers important applications such as real-time uncompressed audio and video.

The RT-VBR traffic class supports time-sensitive services that transmit bits at a variable bit rate (VBR). An example of this traffic class is videoconferencing in which live video frames are compressed and transmitted. For this class of traffic, the ATM network must not introduce significant timing delays since this would distort the live pictures. Thus the delays in cell arrivals and their variances (called jitter) have tight bounds. This class of service is not loss-sensitive, implying that an occasional lost cell does not affect the quality of the picture.

The NRT-VBR traffic class transmits bits at a variable rate but without firm delay bounds. An application suited for NRT-VBR traffic class is multimedia-email. Another application is stored video.

The ABR traffic class is applicable to bit transmission that is bursty. Here, the network services are provided on a best-effort basis. The network provides maximum available throughput with minimum loss. In this service class a feed-back mechanism is in place to control the source behavior. An application suited for ABR traffic class is web-browsing.

In the UBR traffic class, the user sends traffic whenever it wants. An application suited for UBR traffic class is background file transfer. No guarantees or no feedback mechanism is in place. If there is congestion in the network, cells may be dropped. The sources do not reduce traffic during congestion.

2.1.5 AAL Types and QoS

Next we discuss the AAL layers that are in existence today. Currently there are four AALs defined.

- AAL 1: This layer is used for CBR traffic. The Common Part of the Convergence Sublayer (CPCS, Figure 2) performs the following functions: (1) It assembles and disassembles cells, (2) it compensates for delay variation of cells, and (3) it handles lost cells and provides for clock recovery. AAL 1 provides end-to-end timing in the network. It also aligns the application clock to the clock provided by the network.
- AAL 2: This layer is intended for ATM transport of compressed RT-VBR and NRT-VBR audio and video for connection-oriented traffic. A timing relationship between the sender and the receiver is needed. The Motion Picture Experts Group 2 (MPEG-

2) video encoding standard can be used for compression. Furthermore, AAL 2 can be used for bandwidth efficient transport of voice trunking.

- AAL 3/4: This is useful for traffic that is sensitive to loss but not to delay. AAL 3/4 supports both connection-oriented and connectionless VBR traffic. Support for connectionless service is provided at the Service Specific Convergence Sublayer (SSCS) level. AAL 3/4 allows for multiplexing. Multiple sessions from a host can be multiplexed into a single virtual circuit and be transported over AAL 3/4. This layer can operate in message mode or streaming mode. In message mode, the boundaries of messages are preserved. However, in streaming mode, boundaries are not preserved. AAL 3/4 has two kinds of protocol overheads. A message gets a 4 octet overhead, whereas every cell gets an 8 octet overhead. Thus AAL 3/4 is not suitable for short messages.
- AAL 5: Just like AAL 3/4, AAL 5 supports both connection-oriented and connectionless VBR traffic. It also supports both message mode and streaming mode. However, AAL 5 is more efficient than AAL 3/4, since it has no cell overhead. AAL 5 can provide guaranteed delivery with flow control. It can also provide service with no guarantees. Thus AAL 5 is suitable for transport of IP packets.

The ATM traffic classes and the AAL types are closely related to another concept: QoS. It refers to a set of parameters that have been agreed upon between a user and a network to be guaranteed upon connection establishment. These parameters define the quality of service that the network will provide to the user. The agreement is called a traffic contract. As long as the traffic conforms to the agreed upon parameter set, QoS is guaranteed. Some of the important QoS parameters are listed below.

- Peak Cell Rate (PCR): This is the maximum number of cells per second that the user can send. The actual cell rate at any time may be less than the PCR. However under no circumstances may the instantaneous cell rate be greater than the PCR.
- Sustained Cell Rate (SCR): This is the time averaged cell rate the user is generating. For CBR traffic PCR and SCR are exactly equal. For bursty traffic, PCR will generally be much larger than SCR, i.e. $PCR/SCR \gg 1$.
- Minimum Cell Rate (MCR): This is the minimum number of cell per second that the user is generating. If the cell rate in the network is less than the MCR, then the connection is unacceptable. For CBR traffic, MCR is again equal to the PCR. For ABR service, the bandwidth in any instant of time must be bounded by PCR and MCR. For UBR service MCR is set equal to zero.

- **Cell Delay Variation Tolerance (CDVT):** It represents the variation in cell transmission times. If $CDVT = 0$, then a network operating at PCR would deliver 1 cell every $1/PCR$ seconds. This means that the time between the first bit of a cell and the first bit of the next cell is exactly $1/PCR$. For actual sources, this timing is not sustainable. As a result, there is the need for a parameter that takes into account variations in transmission time delays.

Additional QoS parameters that the network provides are given below.

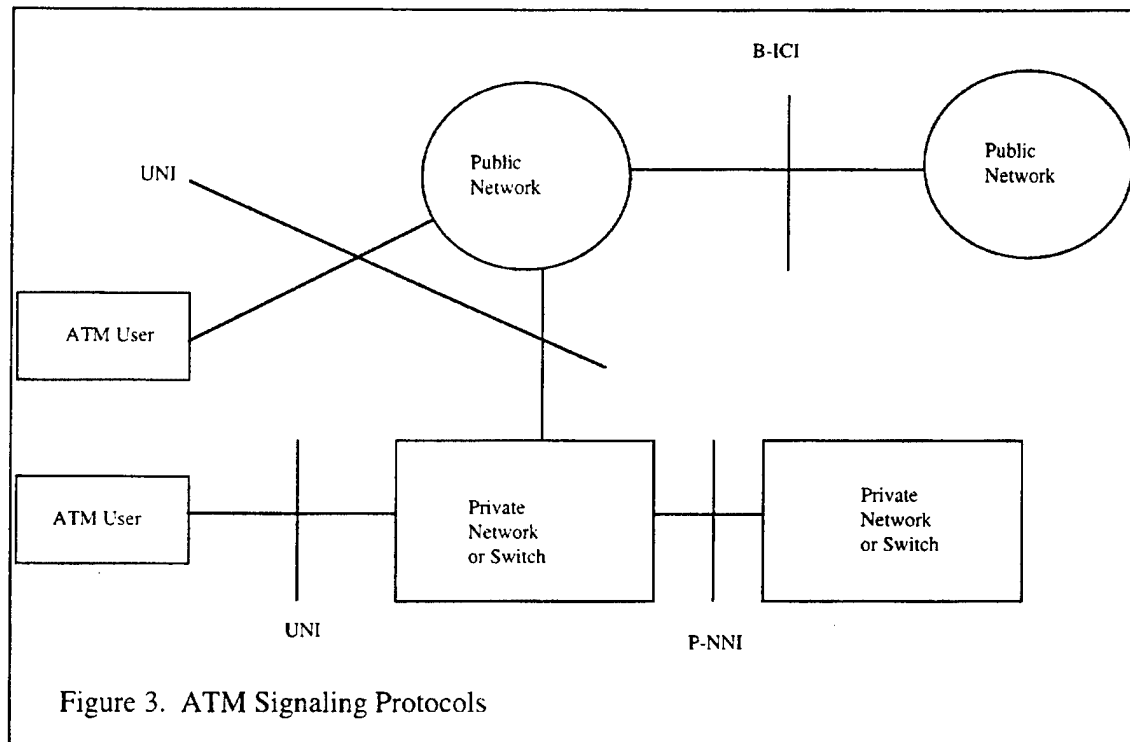
- **Cell Loss Ratio (CLR):** This is the ratio of cells that are lost to the total number of cells that are transmitted.
- **Cell Transfer Delay (CTD):** This is the average time taken between the generation of the first bit of a cell at the source and the receipt of the last bit of a cell at the destination. This delay includes factors such as: coding delay, packetization delay, propagation delay, transmission delay, switching delay, queuing delay, and reassembly delay.
- **Cell Delay Variation (CDV):** This is a measure of how uniformly cells are delivered.
- **Cell Error Ratio (CER):** It is the fraction of cells that are delivered with one or more bits in error. A related quantity is the bit error ratio (BER). It is the ratio of number of bits in error to the total number of bits transmitted. Since fiber technology in ATM networks has made the networks very reliable, BER is an extremely small number.
- **Severely-Errored Cell Block Ratio (SECBR):** This is the fraction of a block of cells that contain an integer number M or more bits in error.
- **Cell Misinsertation Rate (CMR):** It is the number of cells per second that are delivered to a wrong address because of error in destination field of the cell header.

A traffic contract is introduced in the network by specifying some of the parameters mentioned above. The network uses two techniques to adhere to the traffic contract: *traffic shaping* and *traffic policing*. Traffic shaping refers to the process by which the network looks at the source behavior (such as PCR, SCR, and MCR) and buffers the traffic to reduce the PCR. Traffic policing (also called Usage Parameter Control (UPC)) refers to ensuring that sources stay within their negotiated connection setup parameters. The network can take corrective action to punish the violators. Some of these corrective actions include: marking the cells, dropping the cells, delaying the violating cells, and asking a source to correct the violation.

2.1.6 ATM Signaling and Addressing

ATM is a connection-oriented technology. A source that wishes to talk to a destination must establish a connection before data is transmitted. The types of service (e.g., ABR, CBR,

etc.) were discussed in the previous subsection. In this subsection we discuss signaling that is necessary to establish a connection. Signaling refers to a set of procedures and messages used for establishing an end-to-end connection between source and destination stations prior to information transfer. The types of signaling used in ATM networks are discussed below (see Figure 3).



- User-to-Network Interface (UNI) Signaling:

This signaling is used between an ATM end system (ES) and an ATM switch of an ATM network. The current specifications in ATM Forum's UNI signaling is UNI 4.0 specification. The earlier specifications in UNI signaling were UNI 3.0 and UNI 3.1. Most of the available UNI implementations are based on UNI 3.1 and vendors are migrating towards UNI 4.0. To initiate a call, a source ES sends a call Setup message with parameters such as destination address, QoS parameters, and traffic parameters to the ingress ATM switch. This switch invokes a Call Proceeding primitive. The switch invokes a routing protocol to propagate the request across the network to the egress switch. The egress switch then propagates the request to the destination ES. This ES then responds with a Connect message which propagates to the source ES.

along the original path. The source ES receives and acknowledges the Connect message and data transfer takes place.

Some of the features of UNI 3.1 and UNI4.0 signaling are highlighted here. In UNI 3.1 signaling, the root ES can initiate point-to-point and point-to-multipoint connections. The signaling channel is separate from the data transfer channel (in-band signaling). In UNI 3.1 signaling, the network can assign channels for data transfer. UNI 3.1 allows for error recovery and specifies UNI addressing formats.

UNI 4.0 signaling includes several new features to UNI 3.1 signaling. In UNI 4.0, a destination ES can request and join a point-to-multipoint call (leaf-initiated joins). It also supports the capability to negotiate channel numbers for data transfer. It supports QoS classes and also individual parameters for QoS classes. UNI 4.0 also supports anycast signaling (from a source ES to a leader of a group of ESs) and ABR traffic parameters.

- Network-to-Node Interface (NNI) Signaling: This signaling is used between switches or between networks. The signaling between two private switches is called Private NNI (P-NNI), whereas the signaling between two public switches is called Public NNI
- Broadband InterCarrier Interface (B-ICI) Signaling: This signaling is used between two public networks. For example, the signaling between a public switch of a regional carrier and a public switch of a long distance carrier is done using B-ICI.

Along with signaling protocols, the ATM Forum has specified formats of ATM ES addresses. For use within private ATM networks, the Forum has recommended the 20 octet ATM Network Service Access Point (NSAP) address format. For public networks, use of E.164 addresses is recommended. The Forum has also specified the NSAP encoding of E.164 addresses. In the United States, the local authority to assign NSAP addresses is the American National Standards Institute (ANSI) [4]. Within the DoD community, Defense Information Systems Agency (DISA) obtains the NSAP addresses from ANSI and assigns them to users. The NSAP addresses and E.164 addresses contain 6-octet end system identifier addresses which are typically the Ethernet addresses of the network interface cards. An ATM end system registers its NSAP address with a network switch using the Interim Local Management Interface (ILMI). This address registration process then greatly simplifies the establishment of a signaling connection across a UNI.

Up to this point we have discussed the basics of ATM. These basics included key concepts such as ATM traffic classes, AAL types, QoS, signaling and addressing. We next want to ask the question of how to run applications in an ATM network? Since IP-based applications are the ones that are most commonly used, we discuss how IP is run over ATM in the next section.

2.2 IP and ATM

As discussed in the previous section, ATM provides signaling protocols to establish an end-to-end connection. ATM also provides QoS parameters for traffic management purposes. The AAL layers provide varying functionalities to adapt applications to the ATM network. However, to thrive in the market place, ATM must also enable commonly used IP applications to run transparently over ATM. The ATM Forum has actively pursued this area of work to integrate IP traffic into ATM networks. Two of the common methods of running IP in an ATM network are: Classical IP over ATM (put forth by the Internet Engineering Task Force (IETF)) and LAN Emulation (LANE). These are discussed in the following. This is followed by a discussion of MultiProtocol Over ATM (MPOA) which seeks to extend LANE.

2.2.1 Classical IP over ATM

To run connectionless IP datagram protocols over the connection-based ATM network, a scheme for encapsulating IP packets and for mapping IP addresses into ATM addresses is needed. Classical IP running over ATM achieves this in two main steps: packet encapsulation and address resolution. The most common method of packet encapsulation used is Logical Link Control / SubNetwork Access Point (LLC/SNAP) encapsulation [5] across an ATM AAL 5 connection. The resolution of IP addresses into ATM addresses is done using the method prescribed in Request For Comment (RFC) 1577 [6]. An address resolution protocol (ARP) is used to map the IP address either into an NSAP address or an E.164 address. RFC 1577 introduces the notion of a logical IP subnet (LIS) which is a set of IP hosts and routers that connect to a single ATM network and belong to the same IP subnet. Each LIS contains a single ATM ARP server. When a workstation starts up, it is connected to the ATM ARP server which registers its ATM address. If a host does not know the NSAP address of a IP destination, the ATM ARP server resolves that address. Upon receiving the NSAP destination address, the LIS client sets up a connection to that address. After a connection is successfully established to a destination, data is transferred across that connection.

Multicasting is supported according to the prescription given in RFC 2022 [7], which introduces the notion of a multicast address resolution server (MARS). Basically, a set of end points choose to use the MARS to register their membership. A sender has two options for sending multicast data to its members. The first option is to set up a point-to-multipoint VC with the group members as "leaves" and then send data from the sender as the "root". The second option for the source is to send data to a proxy multicast server which then sets up point-to-multipoint connections for transfer of data. FORE system's implementation of classical IP does not support multicasting at the present time. However, multicasting/broadcasting can be done via LANE which we discuss next.

2.2.2 LAN Emulation (LANE)

The second method of running IP applications over ATM is called LANE. It uses ATM as a backbone to interconnect existing LANs. In the LANE specification, there can be multiple logically separate LANs that work transparently as a single LAN over the same ATM network. A single emulated LAN (ELAN) emulates either an Ethernet or a Token Ring. It does not emulate a mixed Ethernet and Token Ring environment. Multiple ELANs may exist on a single ATM network because LANE does not emulate the collision detection algorithm of a legacy LAN.

Integration of LAN technology into ATM presents a challenge because LAN is a connectionless broadcast technology, whereas ATM is a connection-oriented non-broadcast technology. The ATM Forum's response to the challenge was to let every host establish ATM virtual circuits to every other host. This has the potential to proliferate the number of switched virtual circuits to unreasonable numbers.

The LANE v1.0 specification introduces three servers: LAN Emulation Server (LES), Broadcast and Unknown Server (BUS), and the LAN Emulation Configuration Server (LECS). These servers operate as follows:

- LES: The main function of the LES is address look-up. A host sends an Address Resolution Protocol (ARP) request to the LES asking for the IP address that is associated with an ATM address. The LES finds the IP address and sends it to the requesting host. This address is then used to send encapsulated packets to the destination.
- BUS: In a LAN, some applications use broadcasting to find unknown destination addresses. The BUS has virtual connections to all hosts in the emulated LAN. The BUS performs all broadcasting and multicasting and forwards unknown unicast addresses to a default gateway.
- LECS: The LECS entity assigns LANE clients to a particular emulated LAN. LECS provides configuration information and address of the LES.

If there are more than one ELAN, each of these ELANs must have its own LES/BUS entity pair. However, there can be only one LECS for the entire ATM network. Thus the LECS entity can be a single point of failure in the network. The issue of redundant LECSs is addressed in a new version, called LANE v2.0. This enhanced specification provides LANE capabilities such as logical link multiplexing, support for ABR QoS, enhanced multicast support, and support for MPOA.

2.2.3 MPOA

MPOA provides the connectivity of a fully routed environment taking as much advantage of ATM as possible. MPOA specification is a joint effort of the ATM Forum and the IETF. It

provides end-to-end connectivity across an ATM network, for hosts attached directly to the ATM network or indirectly through routers on IP subnetworks.

MPOA is an evolution of LANE. MPOA operates at both Layer 2 (just like a bridge) and at Layer 3 (just like a router). MPOA uses LANE for its Layer 2 forwarding. The key building blocks of MPOA are:

- LANE v2: This protocol is necessary to achieve Layer 2 forwarding.
- Next-Hop Resolution Protocol (NHRP): This is a protocol developed by the IETF and is used to achieve cut-through connection. Two neighboring hosts connected to a large ATM cloud may not realize that they are neighbors because of the way packets are routed. NHRP solves this problem by sending MPOA requests along the routed path to obtain ATM address information which then allows direct ATM virtual connections between MPOA devices.
- MARS: As already mentioned, MARS support the multicast and broadcast needs of Layer 3 protocols. MARS associates Layer 3 multicast group addresses with the ATM interfaces that are group members.

MPOA provides several benefits. It leverages existing infrastructure and allows for flexible changes in the network. Because of MPOA, there is no single point of failure in a wide area network. MPOA compliant products are available now.

2.3 Native ATM Services

As mentioned above, two of the methods of running IP-based applications over ATM are: classical IP over ATM and running IP with the use of LANE. None of the methods (including MPOA) take advantage of the unique benefits of ATM. Either an IP packet or an Ethernet frame has to be created and then overlaid on top of ATM.

Another possible mode of running applications over ATM is the native mode. This refers to an application interfacing with ATM protocols via ATM API syntax defined by the ATM Forum. Running an application in the native mode exploits all the unique features of ATM mentioned earlier.

Currently two ATM APIs have been specified. The Winsock Forum has developed Winsock 2 APIs for Microsoft NT operating system. The Open group has developed the XTI interface for Unix operating systems such as Solaris (Sun Microsystems) and Irix (Silicon Graphics). Both these APIs have been recognized by the ATM Forum as valid mappings of Native ATM Service Access Point (SAP) semantics. FORE systems provides the drivers for both Winsock2 and XTI APIs in their FOREThought (release 5.0 or higher) software for ATM switches.

Several products have been developed that use native ATM. A brief description of two such products are listed below. Both these products are multimedia applications.

- EMMI is a full-duplex Optical Carrier 3, concatenated (OC-3c) multimedia interface product of Lucent Technologies which uses ATM technology to simultaneously transfer audio and video over a LAN or a WAN. EMMI's audio sources consists of two independent stereo channels sampled at 44.1 Kilo Hertz (KHz) by analog-to-digital converter. The sampled information is buffered and processed according to AAL 1 and converted into a bit stream for transmission over the ATM network. The maximum audio bandwidth is 1.85 Mbps. EMMI's video source is a camera or video cassette recorder. The video source is sampled by the video analog-to-digital converter and buffered and encoded using the Joint Photographic Expert Group (JPEG) algorithm. The video information is then processed according to AAL 5 protocol and converted into bit stream for transmission over the ATM network. The video bandwidth has a range from 8 Mbps to 80 Mbps. The EMMI hardware can interface to a host such as a PC, a Macintosh machine or a Sun workstation for generation of data traffic which runs over a TCP/IP stack. The data is converted into the ATM cells using the AAL 5 protocol. The bandwidth usage depends upon the host system configuration. The EMMI interface is intended for applications such as real-time distance learning and tele-medicine.
- CellStack Video and Multi-point Control Unit (MCU) by K-NET Inc. (Plano, Texas), allows several sites across an ATM network to participate in an audio and video session. CellStack supports AAL 5 encapsulation for audio and video streams, and both PVC and SVC signaling. CellStack provides an OC-3c data path for transmitted and received ATM cells. For audio, the maximum bandwidth is 20 Kbps. For video, JPEG compression is used. The CellStack automatically registers the end-system address and generates the 20 octet NSAP address.

LiveLAN 3.1 by PictureTel Corporation is a H.323-based TCP/IP-based video-conferencing application for Personal Computers running Windows 95. PictureTel and FORE Systems are working together to develop and market a LiveLAN product that runs natively over ATM. This product is planned to be released in the first half of 1998.

In summary, we have presented an overview of ATM. We have reviewed ATM traffic classes, the types of ATM adaptation layers, and QoS parameters that the network will provide to the user. ATM signaling interfaces and addressing schemes are discussed. Two methods of running IP applications over ATM are discussed. In classical IP over ATM method specified by the IETF, a mapping of IP addresses into ATM addresses is done. In the LANE method specified by the ATM Forum, every host establishes ATM virtual circuits to every other host. The evolution of LANE into MPOA is also discussed. Finally, native ATM applications are discussed.

Section 3

3. Using ATM to Support ADS for T&E

In this section we present an example of how JADS might be able to utilize ATM. We also examine issues of concern to the JADS T&E community and provide approximate cost estimates.

The JADS ETE test consists of several phases of which phase two is examined in this paper. This phase tests the feasibility of ADS (which includes Distributed Interactive Simulation (DIS) protocol and the evolving Run Time Infrastructure (RTI) framework) to support Developmental T&E (DT&E) and Operational T&E (OT&E) [8]. This test scenario consists of five nodes (Figure 4). The Test, Control, and Analysis Center (TCAC) node is located at Albuquerque, New Mexico and is connected by dedicated T1 links to four other nodes. These nodes are: Fort Hood (III Corps), Texas; White Sands Missile Range (WSMR), White Sands, New Mexico; Grumman Aerospace Laboratory, Melbourne, Florida; and Fort Sill, Oklahoma. The ground station at Grumman is connected to the ground station at Fort Hood by a simulated Surveillance and Control Data Link (SCDL) at a T1 rate. Test control among participating sites is done via one voice channel on the T1 links.

JADS End-to-End Test Phase 2 Topology

1/1/1998


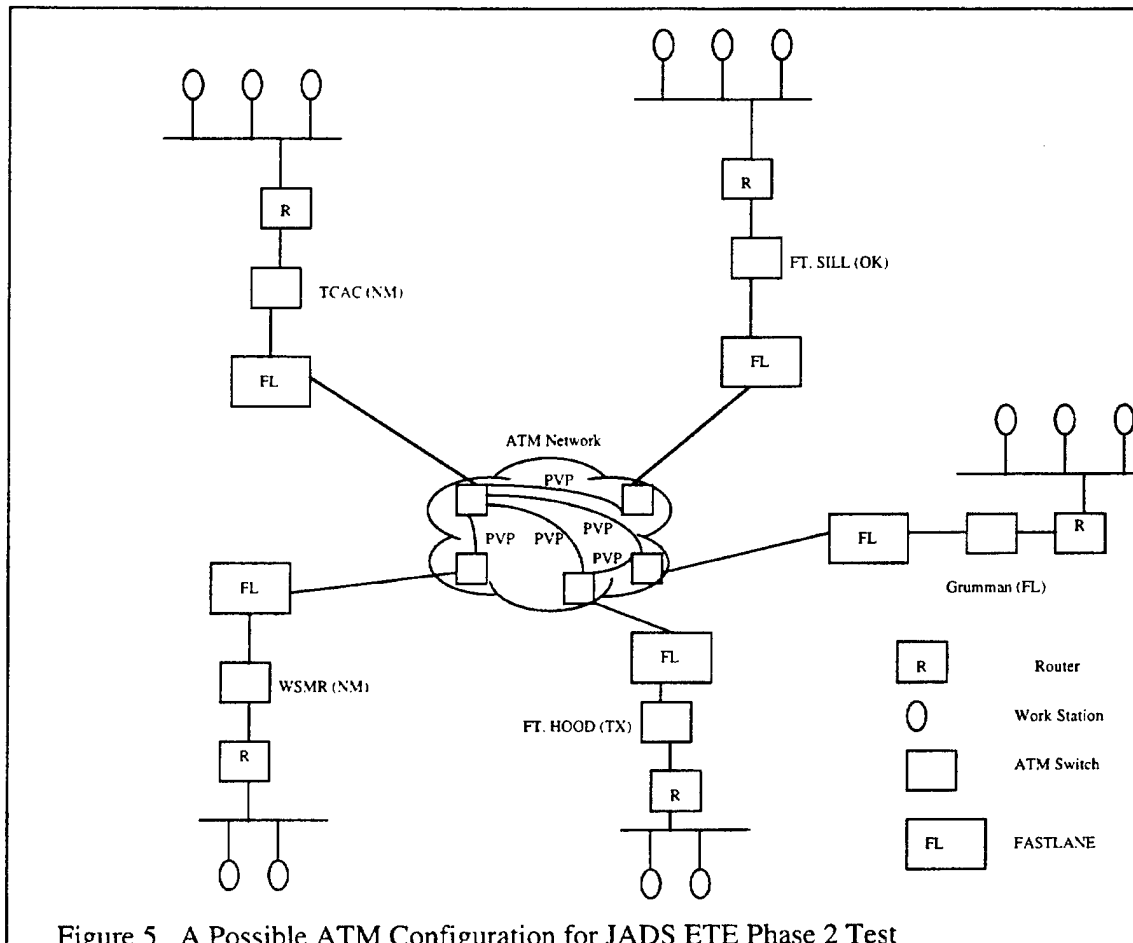


Figure 4. JADS Phase 2 ETE Test Topology

The network baseline requirements for running the JADS ETE phase two test are derived from engineering analysis of expected data traffic. These requirements are (based partly on correspondence with JADS ETE Test Team Lead Lt. Col. Mark McCall):

- The bandwidth (including voice) utilized in any given T1 link does not exceed 775 Kilobits per second (Kbps).
- The aggregate bandwidth exceeds 1.54 Megabits per second (Mbps) but does not exceed 3.08 Mbps.
- The upper limit of latency including allowable errors and radar timelines should be no longer than a second.
- The exercises are classified.

A potential configuration for running a phase 2 ETE test over an ATM network is shown in Figure 5 below.



We have made following assumptions in this configuration:

- Each LAN segment is connected to a router which is connected to an ATM switch (such as FORE ASX-200BX).
- Each ATM switch is connected to a FASTLANE built by GTE Corporation. The FASTLANE is approved by the National Security Agency (NSA) as a type 1 Key Generator-75 (KG-75) encryption device that can operate from Secret level to Top

Secret level. FASTLANE is the replacement for the low speed KIV-7HS bit encryptor that is currently used by JADS. FASTLANE understands ATM protocols and can support up to a total of 4096 cryptographically separated ATM connections. These connections can be combinations of PVCs and SVCs.

- The FASTLANEs are connected to service access points of a public ATM network.
- The five sites have 1.54 Mbps Permanent Virtual Path (PVP) connections as shown in Figure 5. A PVP is a collection of PVCs that can be setup among users.

3.1 JADS ETE Phase Two Test Over ATM

The JADS applications are launched in the workstations shown in Figure 5. The application generates simulation PDUs which are encapsulated in IP packets which communicate over the ATM network via LANE.

The following facts about this set up should be noted:

- Signaling: Connection is initiated between a switch and its user ES via UNI 3.1 signaling. A router and a switch initiate connection via the ILMI mechanism. Connection between a FASTLANE and its adjacent switches is initiated via UNI 3.1 signaling. At the present time FASTLANE doesn't support P-NNI signaling.
- Future FASTLANE release, due in March 1999, is expected to support P-NNI signaling.
- Addressing: As mentioned before, for the site switches, 20 octet NSAP addresses assigned by DISA may be used. The NSAP addresses will include the 6 octet globally unique Ethernet addresses. The emulated LAN interfaces (called asxn interfaces in FORE switches, n being an integer) can be assigned IP addresses as usual.
- LES/BUS and LECS: The LES/BUS and LECS functionality may be chosen to be configured in any site switch, say TCAC.

All the hosts are configured to belong to the same ELAN. After the initial connection set up, traffic flows as if the hosts are all on the same segment.

3.2 Issues and Solutions

ATM is an evolving technology, with specifications still in the process of refinement. It is natural to wonder if ATM network infrastructure is a good investment in terms of capital and technology. Some issues that JADS personnel have raised in this regard were listed earlier. They are again listed below followed by the answers.

- Would the ATM network be public or private? If public, what impact would other traffic have on latency of JADS data?

The JADS community may choose a public ATM service provider to run their simulations. This choice does not impose any performance or security penalties. Data integrity and confidentiality is ensured through FASTLANE encryption devices and the use of PVP pipes.

In the phase two ETE test scenario, latency is not an issue. The ATM network can guarantee much smaller latency over the wide area than the requirement of phase two ETE test. Typically, latencies in the network will be of the order of 100 msec or less. At the time of traffic contract negotiation between the user and the network, stringent delay bounds may be set up. The FASTLANE processes packets at nearly line speed and should not be a problem as far as latencies are concerned.

The ATM network can be part of a DISA Federal network. Some of the other providers have Federal Networks for Government customers, to which JADS can subscribe. The cost estimates are provided in the next section.

- What are the encryption issues? How can encryption be accomplished between sites? Would we have to invest in NES rather than KIV-7?

The low speed KIV-7HSs now used in JADS testing are bit encryptors working at the physical level. These boxes do not understand ATM. The KIV-7HSs should be replaced by FASTLANEs which understand ATM protocol including UNI signaling and transmit cells at speeds very close to the line rate. At the present time, Release 1 (R1) of a FASTLANE box with OC-3c interfaces, costs about \$49 K. Release 3 (R3) of this box is scheduled for March 1999 and will cost \$31 K. This will result in substantial cost reduction and will provide PNNI support. Another product in development by GTE is the TACLANE which encrypts either ATM cells or IP packets, but not both at the same time. TACLANE will provide throughputs of 4 Mbps for IP and about 45 Mbps for ATM. TACLANes are scheduled to be available around December 1998. These units will cost about \$9 K and will be easy to mount on racks.

- What are the bandwidth considerations?

The switches and FASTLANEs have DS-3 ports and OC-3c ports. The bandwidth for DS-3 ports goes up to 45 Mbps and for the OC-3c ports goes up to 155 Mbps. As mentioned earlier, ATM is a scaleable technology. As the bandwidth needs increase, the network shall be able to handle these needs providing QoS guarantees. If the source traffic does not stay within the bounds specified in the traffic contract, cells may be dropped and data will be lost. If the users need more bandwidth, ATM rates become more economical.

- What are the latency considerations? Will latency be fixed or variable?

The ATM network controls traffic through specification of QoS or UPC parameters. One of the UPC parameters is CDVT which is specified in units of microseconds. Thus the network provides for firm delay bounds. For CBR traffic, the latency will consist of fixed transmission delay and fixed propagation delay. For bursty VBR traffic, the aggregate delay will be variable. However, it can be guaranteed to have an upper bound.

- Will small packet size require more processing at simulation sites? Will this add variable latency equal to or greater than WAN savings?

The ATM switches have a non-blocking switch fabric with speed of 2.5 Gbps. With these speeds the processing of cells is very fast. If Cisco routers with AIP cards are used, the fast packet switching option may be turned on. With traditional traffic loads (T1 or its multiples), and switch processing done in the hardware, the Central Processing Unit (CPU) utilization in switches and routers is rather small. The segmentation and reassembly done at the AAL layer will require more processing in the switches. However, with the backplane speeds mentioned above, additional latency is negligible. Even with encryption, the data throughput happens almost at the line rate.

- Will multicast be easier or harder with ATM?

ATM protocols support multicasting. A root node in ATM can set up SVCs to all the leaf nodes of a multicast group. The membership of the group may be dynamic in the sense that the root node may add or drop leaf nodes. For IP-based traffic, the LANE protocol discussed earlier may be used to support multicasting. The establishment of an ELAN is quite straightforward. Initially, all members of the multicasting group need to set up connections to establish an ELAN. As long as the number of hosts in the group stays reasonable, LANE establishment is quite easy. LANE v1.0 does not support selective broadcast. This shortcoming has been removed in LANE v2.0.

- Will ATM be more or less reliable?

In general, experience from Defense Simulation Internet (DSI) and DISA networks shows that ATM networks are pretty reliable. FORE's switch operating system, ForeThought 5.1 implements a distributed LANE environment which is claimed to minimize single point of failure. The network service providers have redundant switches which allow for reconfigurations in case of circuit failure.

- What are cost implications? Will T1 costs go up? Will new boxes be needed at simulation sites?

The recurring monthly charge for a T1 line is assumed to be in the range \$1500 to \$2500 depending on distance. The DISA service rate for ATM CBR traffic with a SCR of 1.5 Mbps is about \$7 K per month. There is also a non-recurring charge of \$1500 for initial site installation. Thus there will be cost increases by going through DISA service. ATM switches and FASTLANE encryption devices will be needed at the communication sites. The costs of these equipment are discussed in the next section.

- What are the cost implications for a private network? What would it take for us to build our own private ATM network (i.e. cost, equipment, training, etc.)?

For T1 and dual T1 ATM pipes, the price difference between a public ATM network and a private ATM networks is expected to be marginal. For higher bandwidth pipes (such as DS3 pipes), and multiple users at a site, there will be substantial savings (up to 50%) in going via a public ATM network [9]. The cost of equipment is discussed in the next section. Because of resource constraints, investigation of training costs and detailed cost implications of private ATM networks could not be made.

In summary, ATM technology is quite capable of meeting the requirements of the JADS community. Since the cost of equipment and services is a major concern, we make a cost analysis in the next section to give a feel for the expenses involved.

3.3 Cost Analysis

In this cost analysis we have picked standard equipment used in modern networks and looked at their costs without attempting to make an exhaustive product comparison.

- The list prices for FORE ASX-200BX switch components are: chassis (switch fabric, 4 empty slots, redundant power supply, and 1 switch control processor) \$16K; a DS-1 module with 4 ports \$5 K; a DS-3 module with 4 ports \$14 K. For Federal Government customers, FORE discounts list prices by 16%. Thus the discounted price for an ASX-200BX switch with a 4 port DS-1 module would cost about \$18 K.
- An ASX-BX200 switch with a 4 port OC-3 module for multimode fiber would cost about \$16.5 K (including discounts).
- Workstations that run the simulation applications may be connected to a router which in turn can be connected to a switch. Alternately, workstations may be equipped with ATM network interface cards (NIC). These cards are available for Peripheral Components Interconnect (PCI) bus or S-bus interfaces. Typically, an ATM NIC card costs less than \$1 K (Efficient Networks is one vendor that manufactures NIC cards).

- For traffic encryption, FASTLANE boxes are used. A FASTLANE with DS-1 or DS-3 support costs about \$47 K and includes the cost of upgrade to Release 3 due in March 1999. Thus the FASTLANEs are a big component of the cost structure. The Release 3 of the FASTLANEs will be cheaper, costing around \$31 K as already mentioned.
- The unofficial DISA service rate for CBR traffic with a SCR of 1.5 Mbps is about \$7 K per month (\$6 K for 1 Mbps) for each pipe. There is also a non-recurring charge of \$1500 for initial site installation. The advantage of ATM becomes transparent if higher bandwidths are needed. At this point the cost savings for the monthly access charges become substantial. For example, a traffic contract for a SCR of 5 Mbps CBR traffic costs only \$8.5 K per month.

Section 4

4. CONCLUSIONS

In this paper we have reviewed some aspects of ATM technology that may be of interest to the JADS community. We have discussed the benefits of ATM internetworking. We have also presented key concepts such as ATM service classes, AAL types, and QoS. We have discussed ATM signaling protocols and the ATM addressing structure. The options for transporting IP traffic over an ATM network have been discussed. The availability of native ATM services for voice and video has also been discussed.

We have considered a particular JADS ETE test scenario and have examined its feasibility of being run over an ATM network. Questions of specific interest to the JADS community, such as latency, bandwidth, and multicasting have been addressed. Finally approximate cost estimates of components and service charges have been given.

From a financial point of view, it would seem that a public ATM network for JADS is not cost-effective as compared to dedicated T1 lines and KIV-7HS bit encryptors in existence now. One big item in the expense column is the high-speed FASTLANE encryptor which replaces KIV-7HS. Another big-ticket item is the ATM switch. As already mentioned, the cost of these boxes is expected to go down in future. Furthermore, alternate boxes may be available in future, which are substantially cheaper and do several levels of encryption. Leaving aside the cost factor, we recommend that ATM network be chosen to carry the JADS simulation data. Advantages of ATM such as QoS guarantees, scalability, fast processing, fast propagation, and firm delay guarantees make ATM very attractive. Ongoing developments in voice and video will make this technology even more attractive in future.

Finally, we mention two areas of further investigation. One area of study is an investigation of running ADS applications directly over ATM. A practical example of how this can be achieved and the benefits that will result is an important area of research. The other area of study is the cost of setting up a private ATM network for JADS. Such a study should include the recurring line charges as a function of bandwidth and the cost of training. It is hoped that as ATM becomes more popular in the future, the capital outlays and service charges become more affordable.

List of References

1. Telco Systems, "Asynchronous Transfer Mode: Bandwidth for the Future", 1992.
2. Tannenbaum, Andrew S., "Computer Networks", Third Edition, Prentice Hall PTR, New Jersey, 1996.
3. McDysan, David E. and Spohn, Darren L., "ATM", McGraw Hill, New York, 1995.
4. Alles, Anthony, "ATM Internetworking", Cisco Systems, 1995.
5. Heinanen, Juha, "Multiprotocol Encapsulation over ATM Adaptation Layer 5", Request For Comment (RFC) 1483, 1993.
6. Laubach, M., "Classical IP and ARP Over ATM", Request For Comment (RFC) 1577, 1994.
7. Armitage, G., "Support for Multicast over UNI 3.0/3.1 based ATM Networks", Request For Comment (RFC) 2022, 1996.
8. Sahu, Devaraj, "A Study of the Defense Simulation Internet (DSI) for the Joint Advance Distributed Simulation (JADS) Project", MITRE Report, April 1998.
9. Yocum, Bernard, Houston Associates, Inc. {703-284-8724}, in telephone conversation, April 14, 1998.

Glossary

AAL	ATM Adaptation Layer
ABR	Available Bit Rate
ADS	Advanced Distributed Simulation
ANSI	American National Standards Institute
API	Application Programming Interface
ARP	Address Resolution Protocol
ATM	Asynchronous Transfer Mode
BER	Bit Error Ratio
B-ICI	Broadband InterCarrier Interface
B-ISDN	Broadband Integrated Services Digital Network
BR	Basic Rate (in ISDN)
BUS	Broadcast and Unknown Server
CBR	Constant Bit Rate
CDV	Cell Delay Variation
CDVT	Cell Delay Variation Tolerance
CER	Cell Error Ratio
CLR	Cell Loss Ratio
CMR	Cell Misinsertion Rate
CPCS	Common Part of Convergence Sublayer (an AAL sublayer)
CPU	Central Processing Unit
CS	Convergence Sublayer (an AAL sublayer)
CTD	Cell Transfer Delay
DIS	Distributed Interactive Simulation
DISA	Defense Information Systems Agency
DS1	Digital Signal Level 1 (1.54 Mbps)
DS3	Digital Signal Level 3 (44.74 Mbps)
DSI	Defense Simulation Internet
DT&E	Developmental Test and Evaluation
ELAN	Emulated LAN
ES	End System
ETE	End-to-end
EW	Electronic Warfare
FY98	Fiscal Year 1998

Gbps	Gigabits per second
HLA	High Level Architecture
IETF	Internet Engineering Task Force
ILMI	Interim Local Management Interface
IP	Internet Protocol
ISDN	Integrated Services Digital Network
ITU	International Telecommunications Union
JADS	Joint Advanced Distributed Simulation
JPEG	Joint Photographic Expert Group
Kbps	Kilobits per second
KG-75	Key Generator 75
KHz	Kilo Hertz
LAN	Local Area Network
LANE	LAN Emulation
LECS	LAN Emulation Configuration Server
LES	LAN Emulation Server
LIS	Logical IP Subnet
LLC	Logical Link Control
MARS	Multicast Address Resolution Server
Mbps	Megabits per second
MCR	Minimum Cell Rate
MCU	Multipoint Control Unit
MPEG-2	Motion Picture Expert Group 2
MPOA	Multiprotocol Over ATM
NHRP	Next Hop Resolution Protocol
NNI	Network-to-node Interface
NRT-VBR	Non Real Time Variable Bit Rate
NSA	National Security Agency
NSAP	Network Service Access Point
OC-3c	Optical Carrier 3, concatenated
OSD	Office of the Secretary of Defense

OSPF	Open Shortest Path First
OT&E	Operational Test and Evaluation
PCI	Peripheral Components Interconnect
PCM	Pulse Code Modulation
PCR	Peak Cell Rate
PDU	Protocol Data Unit
PIM	Protocol Independent Multicasting
PMD	Physical Medium Dependent (a sublayer of ATM physical layer)
P-NNI	Private Network-to-node Interface
PR	Primary Rate (in ISDN)
PVC	Permanent Virtual Circuit
PVP	Permanent Virtual Path
QoS	Quality of Service
RFC	Request for Comment
RT-VBR	Real Time Variable Bit Rate
RTI	Run Time Infrastructure
SAP	Service Access Point
SAR	Segmentation and Reassembly (an AAL sublayer)
SCDL	Surveillance and Control Data Link
SCR	Sustained Cell Rate
SECBR	Severely-Errored Cell Block Ratio
SIT	Systems Integration Test
SNAP	Subnetwork Access Point
SSCS	Service Specific Convergence Sublayer (an AAL sublayer)
SVC	Switched Virtual Circuit
T1	System that transports Digital Signal Level 1 (1.544 Mbps)
TC	Transmission Convergence (a sublayer of ATM physical layer)
TCAC	Test, Control, and Analysis Center
TDM	Time Division Multiplexing
T&E	Test and Evaluation
UBR	Unspecified Bit Rate
UNI	User-to-network Interface
UPC	Usage Parameter Control
VBR	Variable Bit Rate

VCC	Virtual Channel Connection
VCI	Virtual Channel Identifier
VPC	Virtual Path Connection
VPI	Virtual Path Identifier
WAN	Wide Area Network
WSMR	White Sands Missile Range

Distribution List

Internal

Internal

W010

J. S. Quilty

W110

J. C. Slaybaugh

W150

H. J. Carpenter
R. Eftekari

W15D

J. S. Dahmann

W15E

C. E. Walters (5)
E. R. Gonzalez (10)
G. A. Tsoucalas
F. R. Richards
M. Hammond
Files (2)

W15F

M. Adams
S. Chakravorty
D. Sahu (5)
N. Schult

W062

N. J. Slattery
S. Welman

Records Resources (3)

External

Office of the Under Secretary of Defense
(Acquisition)
Deputy Director, Test., systems Engineering &
Evaluation/Systems Assessment
ATTN: Lt. Col. Steven Cameron
The Pentagon, Room 3D1080
Washington, DC 20301-3110 (5)

Colonel Mark Smith
JADS Joint Test Director
JADS JTF
11104 Menaul Blvd., NE
Albuquerque, NM 87112 (10)